EXHIBIT M

Exhibit 5 – SPERT II

’156 Patent

Claim Limitation (Claim 7)	**Exemplary Disclosure**
[**156a**] A device comprising:	SPERT-II discloses a device. Specifically, the SPERT-II is a SBus card. *See, e.g.*: “SPERT-II is a double slot SBus card for use in Sun compatible workstations and is shown in Figure 1. The board contains a T0 vector microprocessor and its memory, a Xilinx FPGA device for interfacing with the host, and various system support devices.” J. Wawrzynek, et al., *SPERT-II: a vector microprocessor system and its application to large problems in backpropagation training*, Proceedings of Fifth International Conference on Microelectronics for Neural Networks, Lausanne, Switzerland, 1996, at 620. 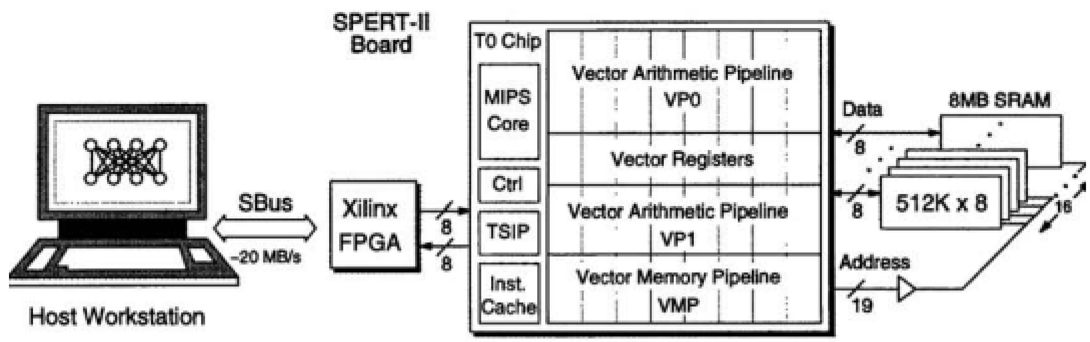 Figure 1: SPERT-II System Organization J. Wawrzynek, et al., *SPERT-II: a vector microprocessor system and its application to large problems in backpropagation training*, Proceedings of Fifth International Conference on Microelectronics for Neural Networks, Lausanne, Switzerland, 1996, at 620. “The T0 processor is a complete single chip implementation of the Torrent architecture. It was fabricated in Hewlett-Packard’s CMOS26B process using 1.0 µm scalable CMOS design rules and two layers of metal. The die measures 16.75mm x16.75mm, and contains 730,701

Exhibit 5 – SPERT II

Claim Limitation (Claim 7)	Exemplary Disclosure
	transistors. T0runs at an internal clock rate of 40MHz.The main components of T0 are the MIPS-II compatible RISC CPU with an on-chip instruction cache, a vector unit coprocessor, a 128-bit wide external memory interface, and an 8-bit wide serial host interface (TSIP) and control unit. The external memory interface supports up to 4 GB of memory over a 128-bit wide data bus. The current SPERT-II board uses 16, 4 Mb SRAM parts to provide 8 MB of main memory." J. Wawrzynek, et al., *SPERT-II: a vector microprocessor system and its application to large problems in backpropagation training*, Proceedings of Fifth International Conference on Microelectronics for Neural Networks, Lausanne, Switzerland, 1996, at 621. "The T0 vector microprocessor integrates an industry-standard MIPS-II 32-bit integer scalar RISC processor with a tightly-coupled fixed-point vector coprocessor. The vector coprocessor contains a central vector register file with 16 vector registers each holding 32 elements of 32 bits, two vector arithmetic units, and a vector memory unit. The two vector arithmetic units each contain 8 parallel pipelines and can produce up to eight 32-bit results per clock cycle. One of the arithmetic units contains 16-bit fixed-point multipliers but otherwise the two pipelines are identical. The memory unit connects to off-chip memory over a 128-bit data bus. T0 has a single flat memory space equally accessible by the scalar unit and any element in the vector unit. T0 is similar in design to vector supercomputers [14], and scalar and vector instructions can be freely intermixed in the single instruction stream. The instruction set was designed to enable object-code compatibility with future higher performance implementations." Krste Asanović, *Programmable Neurocomputing*, The Handbook of Brain Theory and Neural Networks, 2d Ed., 2002, at 3, https://people.eecs.berkeley.edu/~krste/papers/neurocomputing.pdf.

Exhibit 5 – SPERT II

Claim Limitation (Claim 7)	Exemplary Disclosure
	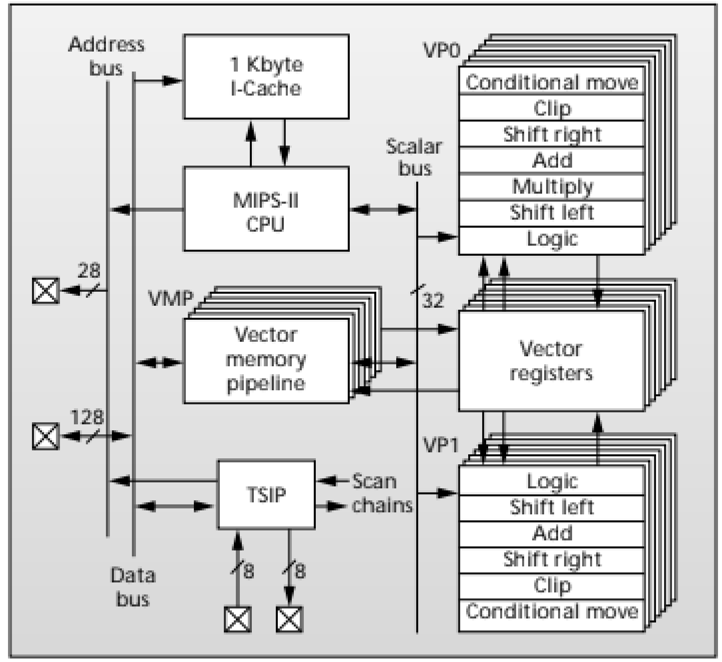 **Figure 1. Block diagram of T0 microarchitecture.** John Wawrzynek, et al., *Spert-II: A Vector Microprocessor System*, COMPUTER, Vol. 29, No. 3, March 1996, at 80.
[**156b**] at least one first low precision high dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value,	SPERT-II discloses at least one first low precision high dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value. *See, e.g.*: "The vector unit contains a vector register file with 16 vector registers, each holding 32 elements of 32 bits each, and three vector functional units, VPO, VP1, and VMP. VPO and VPl are vector arithmetic functional units. With the exception of multiplies, that must execute in VPO, either pipeline can execute any arithmetic operation. The multipliers perform 16-bit x 16-bit multiplies producing 32-bit results. All other arithmetic, logical and shift functions operate on 32 bits. VMP is the vector memory unit, and it handles all vector load/store

Exhibit 5 – SPERT II

Claim Limitation (Claim 7)	Exemplary Disclosure
	operations, scalar load/store operations, and the vector insert/extract operations. All three vector functional units are composed of 8 parallel pipelines, and so can each produce up to 8 results per cycle. The T0 memory interface has a single memory address port, therefore non-unit stride and indexed memory operations are limited to a rate of one element per cycle." J. Wawrzynek, et al., *SPERT-II: a vector microprocessor system and its application to large problems in backpropagation training*, Proceedings of Fifth International Conference on Microelectronics for Neural Networks, Lausanne, Switzerland, 1996, at 621. "The vector unit contains a vector register file and the VP0, VP1, and VMP vector functional units. The vector register file contains 16 vector registers, each holding 32 32-bit elements. VP0 and VP1 are vector arithmetic functional units that can perform 32-bit integer arithmetic and logic operations and support fixed-point scaling, rounding, and saturation. Multiplication is supported only in VP0, with 16-bit × 16-bit multiplies producing 32-bit results. The software performs vector division by long division, using an estimate of the divisor's reciprocal to obtain 12 quotient bits per iteration. VMP, the vector memory unit, handles all vector load/store operations, scalar load/store operations, and vector insert/extract operations. Vectors are addressed in external memory with three types of load/store options—unit stride, nonunit stride, and indexed access. In unit-stride, vector elements occupy consecutive memory locations. In nonunit stride, elements are separated by a constant distance. With indexed access, a vector register provides pointers to the operand vector's elements. This option efficiently implements parallel table-lookup functions for function approximation. It also supports sparse matrix-vector operations. Each vector functional unit is composed of eight parallel pipelines and produces up to eight results per cycle. The T0 memory interface has one memory address port, which limits nonunit stride and indexed memory operations to one element transfer per cycle. A vector register's elements are striped across all eight pipelines. With a maximum vector length of 32, a vector functional unit accepts a new instruction every four cycles. T0 can saturate all three vector functional units by issuing one instruction per cycle to each, leaving the scalar unit an issue slot every four cycles. T0 can thus sustain up to 24 operations per

Exhibit 5 – SPERT II

Claim Limitation (Claim 7)	Exemplary Disclosure
	cycle, a level achieved by several important library routines, such as matrix-vector and matrix-matrix multiplies. All vector pipeline hazards are interlocked in hardware, so instruction scheduling is not required for correctness but may improve performance." John Wawrzynek, et al., *Spert-II: A Vector Microprocessor System*, COMPUTER, Vol. 29, No. 3, March 1996, at 80-81.
[**156c**] wherein the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=5% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.05% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input; and	As explained below and in the Responsive Contentions Regarding Non-Infringement and Invalidity ("Responsive Contentions"), it would have been obvious to one of skill in the art based on the disclosures in SPERT-II (alone or in combination with the reduced precision floating point teachings of Dockser, Tong, Belanovic / Belanovic and Leeser, Lee, Shirazi, Aty, Sudha, and TMS320C32, or the logarithmic format disclosed in GRAPE-3 and Hoefflinger) that the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=5% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.05% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input. SPERT II was designed to perform neural network algorithms at the minimum required precision. "The tendency is increasingly towards general processing elements, more and more closely resembling to RISC or VLIW processors with reduced arithmetic capabilities (no floating-point, reduced precision). Many examples exist, such as the University of Nottingham's *toroidal neural processor*, ICSI's SPERT or the new Philips generation of chips, L-Neuro 2.3. This tendency to generality seems so fundamental as to offset the quest for massive parallelism." Paolo Ienne, *Digital Systems for Neural Networks*, PROC. SPIE 10279, Digital Signal Processing Technology: A Critical Review, April 25, 1995, at 25. Reduced precision computations were commonly used in connection with neurocomputers like the SPERT II. *See, e.g.*:

Exhibit 5 – SPERT II

Claim Limitation (Claim 7)	Exemplary Disclosure
	"Another simplification made possible by ANN algorithms depends on the reduced precision required in most calculation and in most models (Asanović and Morgan 1991; Holt and Baker 1991; Holt and Hwang 1993; Thiran et al. 1994). As a consequence, the designer may avoid area-expensive floating-point arithmetic units and use reduced integer precision (all but one system among those cited in the following sections use fixed-point arithmetic). This results in arithmetic units, registers and data paths using less area on the die and in a reduction of the physical resources devoted to communication (e.g., less IC pins)." Paolo Ienne Lopez, *Programmable VLSI Systolic Processors for Neural Network and Matrix Computations* (unpublished Ph. D. dissertation, École Polytechnique Fédérale de Lausanne), 1996, at 9-10. "Other systems try to take advantage of other peculiarities of ANN algorithms, such as a reduced precision required in the computations. This makes it possible to develop ad-hoc processing elements which are characterized by a small size and cost. In turn, this often enables one to design systems with a very high processing element count and therefore to increase the degree of parallelism." Paolo Ienne, *Digital Systems for Neural Networks*, PROC. SPIE 10279, Digital Signal Processing Technology: A Critical Review, April 25, 1995, at 11. "Limited numeric precision. Many neural algorithms can be coded to require only 8–16 bits of fixed-point arithmetic precision [3]. The reduced precision allows reductions in the area required for arithmetic circuits, particularly multipliers, and also reduces the bandwidth required to transfer operands." Krste Asanović, *Programmable Neurocomputing*, THE HANDBOOK OF BRAIN THEORY AND NEURAL NETWORKS, 2d Ed., 2002, at 3, https://people.eecs.berkeley.edu/~krste/papers/neurocomputing.pdf. Accordingly, SPERT II discloses the use of arithmetic units designed to perform calculations using reduced-precision 16-bit fixed-point math for calculations that typically use 32-bit floating point math, operating on only the 16 most significant bits in the registers. *See, e.g.*: "In the design of the T0 vector microprocessor, the main technique we employ to improve cost-performance over a commercial general purpose processor is to integrate multiple fixed-point datapaths with a high-bandwidth memory system. Fast digital arithmetic units,

Exhibit 5 – SPERT II

Claim Limitation (Claim 7)	Exemplary Disclosure
	multipliers in particular, require chip area proportional to the square of the number of operand bits. In modern microprocessors and digital signal processors a single floating-point unit takes up a significant portion of the chip area. High-precision arithmetic units also requires high memory bandwidth to move large operands. However, for a wide class of problems, full-precision floating-point, or even high-precision fixed-point arithmetic, is not needed. Studies by ourselves and others have shown that for error back-propagation training of neural networks,16-bit weights and 8-bit activation values provide similar training performance to IEEE single-precision floating-point (Asanovic, 1991).” J. Wawrzynek, et al., *SPERT-II: a vector microprocessor system and its application to large problems in backpropagation training*, Proceedings of Fifth International Conference on Microelectronics for Neural Networks, Lausanne, Switzerland, 1996, at 621. “T0 achieves high performance at low cost by integrating multiple fixed-point data paths with a high-bandwidth memory system. Fast digital arithmetic units, such as multipliers and shifters, require chip area proportional to the square of the number of operand bits. In modern micro-processors and digital signal processors, a floating-point unit occupies much of the chip. Higher-precision arithmetic units also need higher memory bandwidth to move large operands. However, many problems do not require full-precision floating-point or even high-precision fixed-point arithmetic. Studies have shown that for error back-propagation neural network training, 16-bit weights and 8-bit activation values perform as well as 32-bit single-precision floating-point.” John Wawrzynek, et al., *Spert-II: A Vector Microprocessor System*, COMPUTER, Vol. 29, No. 3, March 1996, at 80. “Our study shows that, for the speech recognition training problems we have examined, back-propagation training algorithms only require moderate precision. Specifically, 16 b weight values and 8 b output values were sufficient to achieve training and classification results comparable to 32 b floating point. Although these results were gathered for frame classification in continuous speech, we expect that they will extend to many other Connectionist calculations. We have used these results as part of the design of a programmable single chip microprocessor, SPERT. The reduced precision arithmetic permits the use of multiple datapaths per processor. Also, reduced precision operands make more efficient use of valuable processor-memory bandwidth.” John Wawrzynek, et al., *The Design of a Neuro-*

Exhibit 5 – SPERT II

Claim Limitation (Claim 7)	Exemplary Disclosure
	Microprocessor, IEEE TRANSACTIONS ON NEURAL NETWORKS, Vol. 4, No. 3, May 1993, at 394. For the reasons explained in the Responsive Contentions, it would have been obvious to one of skill in the art to have substituted the fixed-point number format used in SPERT-II for a floating-point format that met the claimed minimum range and precision requirements, and to have used the reduced-precision floating-point number formats disclosed in Tong, Dockser, Belanovic / Belanovic and Leeser, Lee Shirazi, Sudha, Aty, and TMS 320C32 or the logarithmic format disclosed in GRAPE-3 and Hoefflinger, either alone or in combination. *See also* Appendix to Responsive Contentions (detailing error rates associated with different mantissa sizes).
[**156d**] at least one first computing device adapted to control the operation of the at least one first LPHDR execution unit;	SPERT-II discloses at least one first computing device adapted to control the operation of the at least one first LPHDR execution unit. *See, e.g.*: "In contrast, T0 integrates an on-chip CPU that acts as controller, and provides fast access to an external memory equally accessible by all datapaths thereby increasing the range of applications that can be run efficiently." J. Wawrzynek, et al., *SPERT-II: a vector microprocessor system and its application to large problems in backpropagation training*, Proceedings of Fifth International Conference on Microelectronics for Neural Networks, Lausanne, Switzerland, 1996, at 622.

Exhibit 5 – SPERT II

Claim Limitation (Claim 7)	Exemplary Disclosure
	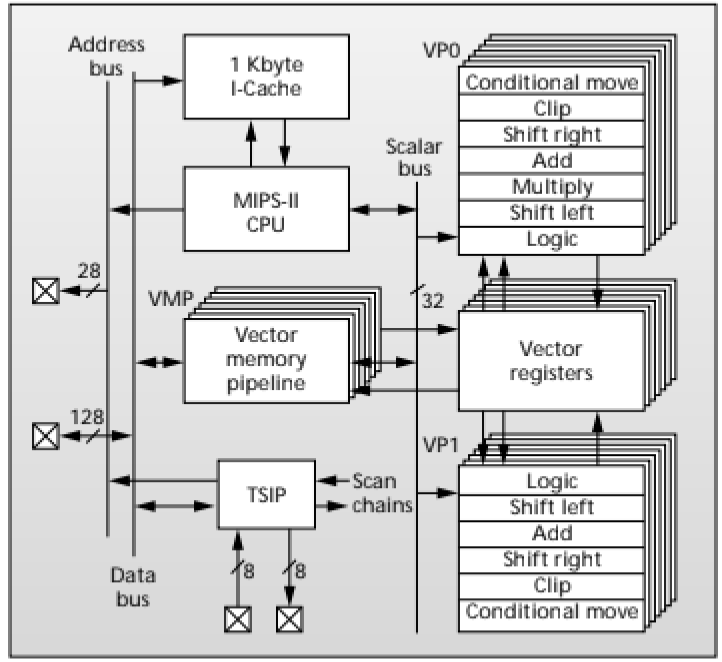 **Figure 1. Block diagram of T0 microarchitecture.** John Wawrzynek, et al., *Spert-II: A Vector Microprocessor System*, COMPUTER, Vol. 29, No. 3, March 1996, at 80.
[**156e**] wherein the at least one first computing device comprises at least one of a central processing unit (CPU), a graphics processing unit (GPU), a field programmable gate array (FPGA), a microcode-based processor, a hardware sequencer, and a state machine;	SPERT-II discloses at least one first computing device comprises at least one of a central processing unit (CPU), a graphics processing unit (GPU), a field programmable gate array (FPGA), a microcode-based processor, a hardware sequencer, and a state machine. *See, e.g.*: "In contrast, T0 integrates an on-chip CPU that acts as controller, and provides fast access to a external memory equally accessible by all datapaths thereby increasing the range of applications that can be run efficiently." J. Wawrzynek, et al., *SPERT-II: a vector microprocessor system and its application to large problems in backpropagation training*, Proceedings of Fifth International Conference on Microelectronics for Neural Networks, Lausanne, Switzerland, 1996, at 622.

Exhibit 5 – SPERT II

Claim Limitation (Claim 7)	Exemplary Disclosure
	"As Figure 1 shows, T0's main components are the MIPS-II-compatible RISC CPU with a 1-Kbyte on-chip instruction cache, a vector unit coprocessor, an external memory inter-face, and an 8-bit-wide serial host interface (TSIP) and control unit. The external memory interface supports up to 4 Gbytes of memory over a128-bit-wide data bus. The system coprocessor provides a 32-bit counter/timer and registers for host synchronization and exception handling." John Wawrzynek, et al., *Spert-II: A Vector Microprocessor System*, COMPUTER, Vol. 29, No. 3, March 1996, at 80. 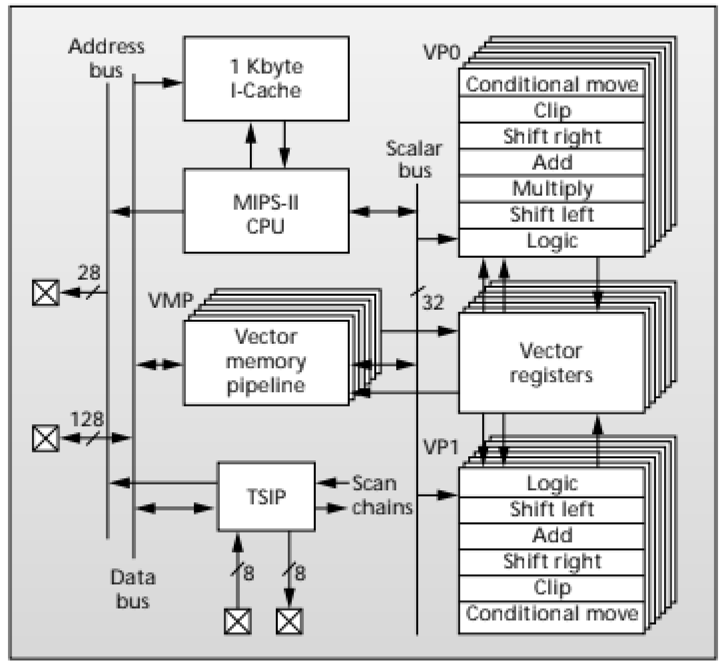 **Figure 1. Block diagram of T0 microarchitecture.** John Wawrzynek, et al., *Spert-II: A Vector Microprocessor System*, COMPUTER, Vol. 29, No. 3, March 1996, at 80.

Exhibit 5 – SPERT II

Claim Limitation (Claim 7)	Exemplary Disclosure
[**156f**] and, wherein the number of LPHDR execution units in the device exceeds by at least one hundred the non-negative integer number of execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide.	SPERT-II teaches or suggests the number of LPHDR execution units in the device exceeds by at least one hundred the non-negative integer number of execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide. The T0 microprocessor includes three vector processing units, each containing eight processing elements. “The vector unit contains a vector register file with 16 vector registers, each holding 32 elements of 32 bits each, and three vector functional units, VPO, VP1, and VMP. VPO and VPl are vector arithmetic functional units. With the exception of multiplies, that must execute in VPO, either pipeline can execute any arithmetic operation. The multipliers perform 16-bit x 16-bit multiplies producing 32-bit results. All other arithmetic, logical and shift functions operate on 32 bits. VMP is the vector memory unit, and it handles all vector load/store operations, scalar load/store operations, and the vector insert/extract operations. All three vector functional units are composed of 8 parallel pipelines, and so can each produce up to 8 results per cycle. The T0 memory interface has a single memory address port, therefore non-unit stride and indexed memory operations are limited to a rate of one element per cycle.” J. Wawrzynek, et al., *SPERT-II: a vector microprocessor system and its application to large problems in backpropagation training*, Proceedings of Fifth International Conference on Microelectronics for Neural Networks, Lausanne, Switzerland, 1996, at 621. “The vector unit contains a vector register file and the VP0, VP1, and VMP vector functional units. The vector register file contains 16 vector registers, each holding 32 32-bit elements. VP0 and VP1 are vector arithmetic functional units that can perform 32-bit integer arithmetic and logic operations and support fixed-point scaling, rounding, and saturation. Multiplication is supported only in VP0, with 16-bit × 16-bit multiplies producing 32-bit results. The software performs vector division by long division, using an estimate of the divisor’s reciprocal to obtain 12 quotient bits per iteration. VMP, the vector memory unit, handles all vector load/store operations, scalar load/store operations, and vector insert/extract operations. Vectors are addressed in external memory with three types of load/store options—unit stride,

Exhibit 5 – SPERT II

Claim Limitation (Claim 7)	Exemplary Disclosure
	nonunit stride, and indexed access. In unit-stride, vector elements occupy consecutive memory locations. In nonunit stride, elements are separated by a constant distance. With indexed access, a vector register provides pointers to the operand vector's elements. This option efficiently implements parallel table-lookup functions for function approximation. It also supports sparse matrix-vector operations. Each vector functional unit is composed of eight parallel pipelines and produces up to eight results per cycle. The T0 memory interface has one memory address port, which limits nonunit stride and indexed memory operations to one element transfer per cycle. A vector register's elements are striped across all eight pipelines. With a maximum vector length of 32, a vector functional unit accepts a new instruction every four cycles. T0 can saturate all three vector functional units by issuing one instruction per cycle to each, leaving the scalar unit an issue slot every four cycles. T0 can thus sustain up to 24 operations per cycle, a level achieved by several important library routines, such as matrix-vector and matrix-matrix multiplies. All vector pipeline hazards are interlocked in hardware, so instruction scheduling is not required for correctness but may improve performance." John Wawrzynek, et al., *Spert-II: A Vector Microprocessor System*, COMPUTER, Vol. 29, No. 3, March 1996, at 80-81. SPERT-II is modular, so additional microprocessors can be added to achieve higher data parallelism. "An important goal in the SPERT design is to prototype ideas for a parallel processing node that will be used in a future, scalable, MIMD multiprocessor system. Such a system would target large, irregular network structures." John Wawrzynek, et al., *The Design of a Neuro-Microprocessor*, IEEE TRANSACTIONS ON NEURAL NETWORKS, Vol. 4, No. 3, May 1993, at 398. For the reasons explained above and in the Responsive Contentions, it would have been obvious to one of skill in the art based on the disclosures in SPERT-II (alone or in combination with the parallel processing elements teachings of MacMillan and Cloutier) that the number of

Exhibit 5 – SPERT II

Claim Limitation (Claim 7)	Exemplary Disclosure
	LPHDR execution units in the device exceeds by at least one hundred the non-negative integer number of execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide.

Exhibit 5 – SPERT II

’273 Patent

Claim Limitation (Claim 53)	Exemplary Disclosure
[**273a**] A device:	SPERT-II discloses a device. Specifically, the SPERT-II is a SBus card. *See* [**156a**].
[**273b**] comprising at least one first low precision high-dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value,	SPERT-II discloses at least one first low precision high dynamic range (LPHDR) execution unit adapted to execute a first input signal representing a first numerical value to produce a first output signal representing a second numerical value. *See* [**156b**].
[**273c**] wherein the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=5% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X % of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.05% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input;	For the reasons explained above and in the Responsive Contentions, it would have been obvious to one of skill in the art based on the disclosures in SPERT-II (alone or in combination with the reduced precision floating point teachings of Dockser, Tong, Belanovic / Belanovic and Leeser, Lee, Shirazi, Aty, Sudha, and TMS320C32 or the logarithmic format disclosed in GRAPE-3 and Hoefflinger) that the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=5% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.05% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input. *See* [**156c**]; *see also* Appendix to Responsive Contentions (detailing error rates associated with different mantissa sizes).
[**273d**] wherein the number of LPHDR execution units in the device exceeds by at least one hundred the non-negative integer number of execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide.	SPERT-II teaches or suggests the number of LPHDR execution units in the device exceeds by at least one hundred the non-negative integer number of execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide. *See* [**156f**].

Exhibit 5 – SPERT II

’961 Patent

Claim Limitation (Claim 4)	Exemplary Disclosure
[**961a**] A device comprising:	SPERT-II discloses a device. Specifically, the SPERT-II is a SBus card. *See* [**156a**].
[**961b**] at least one first low precision high-dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value,	SPERT-II discloses at least one first low precision high dynamic range (LPHDR) execution unit adapted to execute a first input signal representing a first numerical value to produce a first output signal representing a second numerical value. *See* [**156b**].
[**961c**] wherein the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=10% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.2% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input; and	For the reasons explained above and in the Responsive Contentions, it would have been obvious to one of skill in the art based on the disclosures in SPERT-II (alone or in combination with the reduced precision floating point teachings of Dockser, Tong, Belanovic / Belanovic and Leeser, Lee, Shirazi, Aty, Sudha, and TMS320C32 or the logarithmic format disclosed in GRAPE-3 and Hoefflinger) that the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=10% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.2% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input. *See* [**156c**]; *see also* Appendix to Responsive Contentions (detailing error rates associated with different mantissa sizes).
[**961d**] at least one first computing device adapted to control the operation of the at least one first LPHDR execution unit.	SPERT-II discloses at least one first computing device adapted to control the operation of the at least one first LPHDR execution unit. *See* [**156d**].

Exhibit 5 – SPERT II

Claim Limitation (Claim 13)	Exemplary Disclosure
[**961e**] A device comprising:	SPERT-II discloses a device. Specifically, the SPERT-II is a SBus card. *See* [**156a**].
[**961f**] a plurality of components comprising:	*See* SPERT-II discloses at least one first low precision high dynamic range (LPHDR) execution unit adapted to execute a first input signal representing a first numerical value to produce a first output signal representing a second numerical value. *See* **[156b].** *See also* SPERT-II discloses at least one first computing device adapted to control the operation of the at least one first LPHDR execution unit. *See above* [**156d**].
[**961g**] at least one first low precision high-dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value,	SPERT-II discloses at least one first low precision high dynamic range (LPHDR) execution unit adapted to execute a first input signal representing a first numerical value to produce a first output signal representing a second numerical value. *See* [**156b**].
[**961h**] wherein the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=10% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.2% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input.	For the reasons explained above and in the Responsive Contentions, it would have been obvious to one of skill in the art based on the disclosures in SPERT-II (alone or in combination with the reduced precision floating point teachings of Dockser, Tong, Belanovic / Belanovic and Leeser, Lee, Shirazi, Aty, Sudha, and TMS320C32 or the logarithmic format disclosed in GRAPE-3 and Hoefflinger) that the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=10% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.2% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input. *See* [**156c**]; *see also* Appendix to Responsive Contentions (detailing error rates associated with different mantissa sizes).